

Efficient Associative Search in Brain-Inspired Hyperdimensional Computing

Mohsen Imani, Justin Morris, Helen Shu, Shou Li, and Tajana Rosing

University of California at San Diego

Editor's note:

This article describes a method for efficient hypervector operations using a grouping strategy for reduced computations. Quantization is used for reducing the number of multiplications, whereas caching of magnitude is used for eliminating redundant computations.

—Shao-Wen Yang, Intel Corporation

■ **THE EMERGENCE OF** the Internet of Things has created an abundance of small embedded devices. Many of these devices may be used for cognitive tasks such as face detection, speech recognition, image classification, activity monitoring, and so on. Learning algorithms such as deep neural networks (DNNs) give accurate results for cognitive tasks [1]. However, these embedded devices have limited resources and cannot run such resource-intensive algorithms. Instead, many of them send the data they collect to the cloud server for feature analysis. However, this is not desirable due to privacy concerns, security concerns, and network resources. Thus, we need more efficient lightweight classifiers in order to perform such cognitive tasks on the embedded systems.

Brain-inspired hyperdimensional (HD) computing can be used as a lightweight classifier to perform cognitive tasks on resource-limited systems [2]. HD

computing is modeled after how the brain works, using patterns of neural activity instead of computational arithmetic. Past research utilized high-dimension vectors ($D \geq 10,000$) called *hypervectors* to represent neural patterns. It showed that HD computing is capable of providing high-accuracy results for a variety of

tasks such as language recognition, face detection, speech recognition, classification of time-series signals, and clustering [3]–[7]. Results are obtained at a much lower computational cost as compared to other learning algorithms.

HD computing performs the classification task after encoding all data points to high-dimensional space. The HD training happens by linearly combining the encoded hypervectors and creating a hypervector representing each class. In inference, HD uses the same encoding module to map a test data point to high-dimensional space. Then the classification task checks the similarity of the encoded test hypervector with all pretrained class hypervectors. This similarity check is the main HD computation during the inference; often done with a cosine, which involves a large number of costly multiplications, grows with the number of classes [8]. Given an application with k classes, inference requires $k \times D$ additions and multiplications to perform, where D is the hypervector dimension. Thus, this similarity check can be costly for embedded devices with limited resources.

Digital Object Identifier 10.1109/MDAT.2019.2919954

Date of publication: 30 May 2019; date of current version: 6 February 2020.

In this article, we propose a robust and efficient solution to reduce the computational complexity and cost of HD computing while maintaining comparable accuracy. The proposed HD framework exploits the mathematics in high-dimensional space in order to limit the number of classes checked upon inference, thus reducing the number of computations needed for query requests. We add a new layer before the primary HD layer to decide which subset of class hypervectors should be checked as possible classes for the output class. This reduces the number of additions and multiplications needed for inference. In addition, our framework removes the costly multiplication from the similarity check by quantizing the values in the trained HD model with the power of two values. Our approach integrates quantization with the training process in order to adapt the HD model to work with the quantized values. We have evaluated our proposed approach on three practical classification problems. Our evaluations show that the proposed design is 11.6× more energy efficient and 8.3× faster as compared to the baseline HD, while providing similar classification accuracy.

Hyperdimensional computing

HD computing uses long vectors with dimensionality in the thousands [2]. There are many nearly orthogonal vectors in high-dimensional space. HD combines these hypervectors with well-defined vector operations while preserving most of their information. No component has more responsibility to store any piece of information than any other component because hypervectors are holographic and (pseudo) random with independent and identically distributed (i.i.d.) components and full holistic representation. The mathematics governing the high-dimensional space computations enables HD to be easily applied to many different learning problems.

Figure 1 shows an overview of the structure of the HD model. HD consists of an encoder, trainer, and associative search block. The encoder maps data points into high-dimensional space. These hypervectors are then combined in a trainer block to form class hypervectors, which are then stored in an associative search block. In inference, an input test data is encoded to high-dimensional space using the same encoder as the training module. The classifier uses cosine similarity to check the similarity of the

encoded hypervector with all class hypervectors and find the most similar one.

Encoding

Consider a feature vector $\mathbf{v} = \langle v_1, \dots, v_n \rangle$. The encoding module takes this n -dimensional vector and converts it into a D -dimensional hypervector ($D \gg n$). The encoding is performed in three steps, which we describe in the following.

We use a set of precomputed level or base hypervectors to consider the impact of each feature value [9]. To create these level hypervectors, we compute the minimum and maximum feature values among all data points, \mathbf{v}_{\min} and \mathbf{v}_{\max} , then quantize the range of $[\mathbf{v}_{\min}, \mathbf{v}_{\max}]$ into Q levels, $\mathbf{L} = \{\mathbf{L}_1, \dots, \mathbf{L}_Q\}$. Each of these quantized scalars corresponds to a D -dimensional hypervector [9].

Once the base hypervectors are generated, each of the n elements of the vector v is independently quantized and mapped to one of the base hypervectors. The result of this step is n different binary hypervectors, each of which is D -dimensional. In the last step, the n (binary) hypervectors are combined into a single D -dimensional (nonbinary) hypervector. To differentiate the impact of each feature index, we devise ID hypervectors, $\{\mathbf{ID}_1, \dots, \mathbf{ID}_n\}$. An ID hypervector has the binarized dimensions, i.e., $\mathbf{ID}_i \in \{0, 1\}^D$. We create ID s with random binary values so that the ID hypervectors of different feature indexes are nearly orthogonal:

$$\delta(\mathbf{ID}_i, \mathbf{ID}_j) \simeq D/2 \quad (i \neq j \text{ \& } 0 < i, j \leq n)$$

where the similarity metric, $\delta(\mathbf{ID}_i, \mathbf{ID}_j)$, is the Hamming distance between the two ID hypervectors.

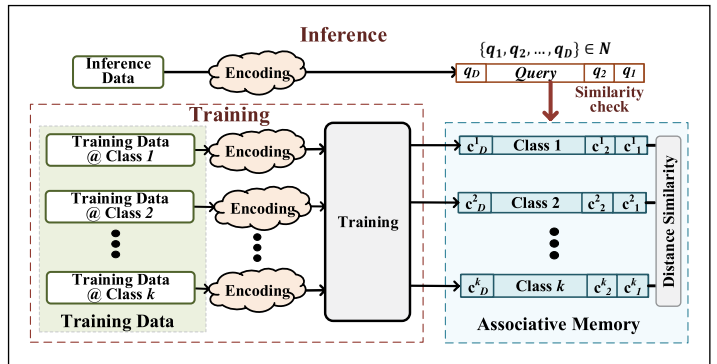


Figure 1. Overview of HD computing performing the classification task.

The orthogonality of ID hypervectors is ensured as long as the hypervector dimensionality is large enough compared to the number of features in the original data point ($D \gg n$). As Figure 1 shows, the aggregation of the n binary hypervectors is computed as follows:

$$\mathbf{H} = \mathbf{ID}_1 \oplus \mathbf{L}_1 + \mathbf{ID}_2 \oplus \mathbf{L}_2 + \dots + \mathbf{ID}_n \oplus \mathbf{L}_n$$

where \oplus is XOR operation, \mathbf{H} is the aggregation, and \mathbf{L}_i is the binary hypervector corresponding to the i^{th} feature of vector \mathbf{v} .

Classification in HD space

In HD, training is performed in high-dimensional space by elementwise addition of all encoded hypervectors in each existing class. The result of training will be k hypervectors with D dimensions, where k is the number of classes. For example, the i^{th} class hypervector can be computed as: $\mathbf{C}^i = \sum_{\mathbf{v} \in \text{class}_i} \mathbf{H}_{\mathbf{v}}$.

In inference, HD uses encoding and associative search for classification. First, HD uses the same encoding module as the training module to map a test data point to a query hypervector. In HD space, the classification task then checks the similarity of the query with all class hypervectors. The class with the highest similarity to the query is selected as the output class. Since in HD, information is stored as the pattern of values, the cosine is a suitable metric for similarity check.

HD computing challenges

To understand the main bottleneck of HD computing during inference, we evaluate the HD

inference on three practical classification applications, including speech recognition [10], activity recognition [11], and image recognition [12]. All evaluations are performed on Intel i7 7600 CPU with 16-GB memory. Our results show that the associative search takes about 83% of the total inference execution. This is because the cosine similarity has many multiplications between a query and the class hypervectors. Prior work tried to binarize the HD model after the training [3]. This method simplifies the cosine similarity to Hamming distance measurement, which can run faster and more efficiently in hardware. However, our evaluation of three practical applications shows that binarization reduces the classification accuracy of HD with integer model by 11.4% on an average. This large accuracy drop forces HD to use cosine similarity which has a significant cost when running on embedded devices.

HD computing acceleration

In this article, we propose three optimization methods that reduce the cost of associative search during inference by at least an order of magnitude. Figure 2 shows an overview of the proposed optimizations. The first approach simplifies the cosine similarity calculations to dot products between the query and class hypervectors. The second reduces the number of required operations in the associative search by adding a category layer to HD computing which decides what subset of class hypervectors needs to be checked for the output class. The third removes the costly multiplications from the similarity

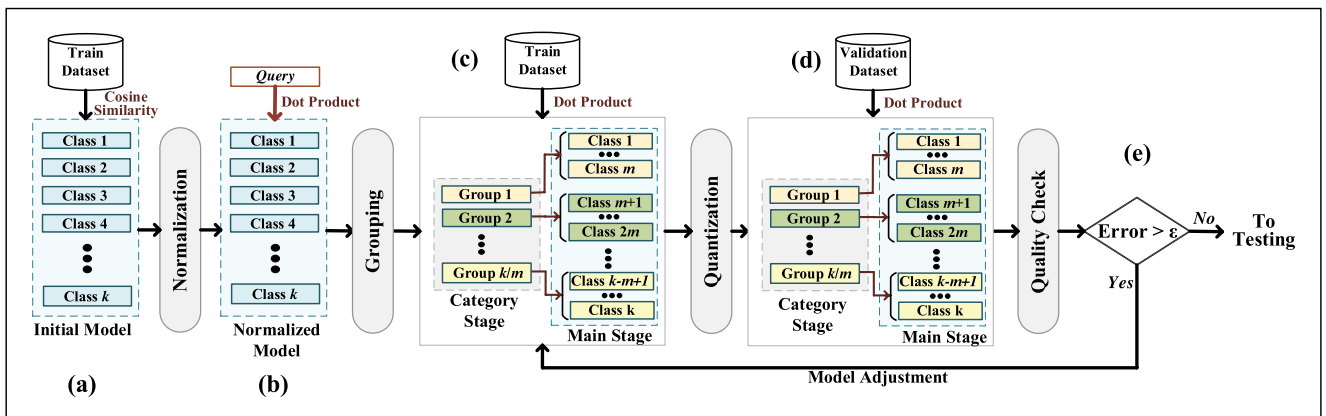


Figure 2. Overview of proposed optimization approaches to improve the efficiency of associative search.

check by quantizing the HD model after training. In the following subsections, we explain the details of each proposed approach.

Similarity check: Cosine or dot product?

During inference, HD computation encodes input data to a query hypervector, $\mathbf{H} = \{h_D, \dots, h_2, h_1\}$. Associative memory then measures the cosine similarity of this query with k stored class hypervectors $\{\mathbf{C}^1, \dots, \mathbf{C}^k\}$, where $\mathbf{C}^i = \{c_D^i, \dots, c_2^i, c_1^i\}$ is the class hypervector corresponding to the i^{th} class (Figure 2a). The cosine similarity can be expressed as $\delta = \mathbf{H} \cdot \mathbf{C}^i / |\mathbf{H}| |\mathbf{C}^i|$, where $\mathbf{H} \cdot \mathbf{C}^i$ indicates the dot product between the hypervectors, and $|\mathbf{H}| = \sqrt{\mathbf{H} \cdot \mathbf{H}}$ and $|\mathbf{C}^i| = \sqrt{\mathbf{C}^i \cdot \mathbf{C}^i}$ show the magnitudes of the query and class hypervector. However, it is very expensive to calculate the operand magnitude every time. During the similarity check, the query hypervector is common between all classes. Thus, we can skip the calculation of the query magnitude, since the goal of HD is to find the maximum relative similarity, and not the exact cosine values. On the other hand, as shown in Figure 2b, the magnitude of each class hypervector can be computed once after the training. Therefore, the associative search can store the normalized class hypervectors $(\mathbf{C}^i / |\mathbf{C}^i| \text{ for } i \in 1, \dots, k)$. This speeds up the similarity at inference by about $3\times$ as compared to cosine.

Two level search

Although the dot product reduces the cost of the associative search, this similarity check still involves many computations. For example, for an application with k classes, associative search computes $k \times D$ multiplication/addition operations, where D is the hypervector dimension. In addition, in existing HD computing approaches [3], [9], the cost of associative search increases linearly with the number of classes. For example, speech recognition with $k = 26$ classes has $13\times$ more computations and $5.2\times$ slower inference as compared to face detection with $k = 2$ classes. Since embedded devices often do not have enough memory and computing resources, processing HD applications with large numbers of classes is more inefficient.

We propose a method that has a two-layer classification: category and main stages. Figure 2c shows an overview of the proposed approach. First, we group the trained class hypervectors into

k/m categories based on their similarity, where k and m are the number of classes and group size, respectively. For example, $m = 2$ indicates that we group every two class hypervectors into a single hypervector. Next, we build a new HD model, called *category stage*, which stores all k/m group hypervectors. Instead of searching k hypervectors to classify a data point, we first search the category stage to identify the group of classes that the query belongs to (among k/m group hypervectors). Afterward, we continue the search in the main HD stage, but only with the class hypervectors corresponding to the selected group.

Quantization

Although grouping the class hypervectors reduces the number of computations, the dot product similarity check still has many costly multiplications. In this work, we propose a method which removes the majority of the multiplications from the HD similarity check. After training the HD model, we quantize each class value to the closest power of two ($2^i, i \in \mathbb{Z}$). This eliminates the multiplication by allowing bit shift operations. However, this quantization is not error free. For example, the closest power of two for the number 46 would be either 32 or 64. Both of these numbers are far from the actual value. This approximation can add large errors to HD computing. The amount of error depends on the number of classes and how similar the classes are. For applications with many classes or highly correlated class hypervectors, this quantization can have a large impact on the classification accuracy.

In this article, we look at the possibility of performing a more precise but lower power quantization approach (Figure 2d). Our method assigns each class element to a combination of two power of two values ($2^i + 2^j, i \& j \in \mathbb{Z}$). This quantization can assign each trained class element to a value which is much closer to the actual element. For example, using this approach, the number 46 can be assigned to $48 = 2^5 + 2^4$, which is very close to the actual value. This enables more precise quantization with correspondingly lower impact on accuracy. This strategy implements multiplication using two shifts and a single add operation, which is still faster and more efficient than the actual multiplication. After training the HD model, we assign the class elements in both category and main stages to the closest quantized value.

Since class hypervectors are highly correlated, even this quantization can have a large impact on the classification accuracy. Quantization introduces this quality loss because the HD model is not trained to work with the quantized values. In order to ensure a minimum quality loss, we integrate quantization with the HD model retraining. This enables the HD model to learn how to work with quantized values. In the algorithm explained in the “Training and inference in the proposed design” section, after getting a new adjusted model, we quantize all hypervector values in the category and main stages. This approach reduces the possible quality loss due to quantization. In the “Accuracy” section, we discuss the impact of quantization on HD classification accuracy.

Training and inference in the proposed design

Training: Figure 2 shows the training process of HD with grouped hypervectors. We first train a normal HD computing model, where each hypervector represents an existing class (Figure 2a). Next, we normalize the class hypervectors (Figure 2b) and then check the similarity of the trained class hypervectors in order to group the classes. In our approach, we group every m class hypervector into a single hypervector, and as a result, we have k/m group hypervectors (Figure 2c). The grouping is done by checking the similarity of class hypervectors in pairs and classes were merged with the highest similarity. The selected class hypervectors are added together to generate a group hypervector. We store these k/m group hypervectors into the category stage. Then, we quantize the values of the grouped model (Figure 2d). This one-shot trained model can be used to perform the classification task at inference.

Model adjustment: To get better classification accuracy, we can adjust the HD model with the training data set for a few iterations (Figure 2e). The model adjustment starts in the main HD stage. During a single iteration, HD checks the similarity of all training data points, say \mathbf{H} , with the current HD model. If data are wrongly classified by the model, HD updates the model by 1) adding the data hypervector to a class to which it belongs and 2) subtracting it from a class with which it was wrongly matched.

$$\text{Main} \begin{cases} \tilde{\mathbf{C}}_{\text{main}}^c = \mathbf{C}_{\text{main}}^c + \mathbf{H}, \text{ where } \mathbf{C}_{\text{main}}^c \text{ is correct} \\ \tilde{\mathbf{C}}_{\text{main}}^w = \mathbf{C}_{\text{main}}^w - \mathbf{H}, \text{ where } \mathbf{C}_{\text{main}}^w \text{ is wrong} \end{cases}$$

We similarly update two corresponding hypervectors in the category stage by adding and subtracting the query hypervector.

$$\text{Category} \begin{cases} \tilde{\mathbf{C}}_{\text{category}}^c = \mathbf{C}_{\text{category}}^c + \mathbf{H}, \text{ where } \mathbf{C}_{\text{main}}^c \in \mathbf{C}_{\text{category}}^c \\ \tilde{\mathbf{C}}_{\text{category}}^w = \mathbf{C}_{\text{category}}^w - \mathbf{H}, \text{ where } \mathbf{C}_{\text{main}}^w \in \mathbf{C}_{\text{category}}^w \end{cases}$$

The model adjustment needs to be continued for a few iterations until the HD accuracy stabilizes over the validation data, which is a part of the training data set. After training and adjusting the model offline, it can be loaded onto embedded devices to be used for inference.

Inference: The proposed approach works very similarly to the baseline HD computing, except there are two stages now. First, we check the similarity of a query hypervector in the category stage. A category hypervector with the highest cosine similarity is selected to continue the search in the main stage.

Here, we check the similarity of the query hypervector against the classes within the selected category. For example, in Figure 2c, if group 2 had the highest cosine similarity with the query hypervector, then only the green class hypervectors are selected for search in the main stage. Finally, a class with the highest cosine similarity in the main stage is selected as the output class. This approach reduces the number of required operations. For an application with k classes, our approach reduces the number of required similarity checks from k to $k/m + m$ hypervectors. For example, for an application with $k = 16$ and $m = 4$, the number of required operations is reduced by a factor of 2.

Evaluation

Experimental setup

We perform HD training and retraining using C++ implementation on Intel Core i7 processor with 16-GB memory (4 core, 2.8 GHz). We describe the inference functionality using RTL System-Verilog and use standard digital ASIC flow to design dedicated hardware. For the synthesis, we use Synopsys Design Compiler with the Taiwan Semiconductor Manufacturing Company (TSMC) 45-nm technology library and the general purpose processor with high- V_{TH} cells. We extract the design switching activity using ModelSim, and measure the power consumption of HD designs using Synopsys Prime-Time at the (1 V, 25 °C, TT) corner.

Table 1. Classification accuracy of HD with integer (baseline) and quantized model.

# of Classes in a Group	Integer Model				Quantized (2^i)				Quantized ($2^i + 2^j$)			
	1	2	3	4	1	2	3	4	1	2	3	4
Speech Recognition	91.72%	88.39%	88.01%	87.04%	91.73%	83.00%	85.70%	83.39%	90.64%	88.45%	88.26%	88.39%
Activity Recognition	95.50%	95.49%	95.37%	94.98%	93.50%	92.73%	94.92%	94.72%	94.27%	94.34%	95.62%	94.72%
Image Recognition	92.97%	93.25%	91.02%	89.11%	89.99%	92.54%	91.05%	80.62%	91.51%	93.00%	92.28%	89.31%

We test the efficiency of the proposed approach on the following three practical applications:

Speech Recognition (ISOLET) [10]: Recognize voice audio of the 26 letters of the English alphabet. The training and testing data sets are taken from the Isolet data set. This data set consists of 150 subjects speaking each letter of the alphabet twice. The speakers are grouped into sets of 30 speakers. The training of hypervectors is performed on Isolet 1, 2, 3, 4, and tested on Isolet 5.

Activity Recognition (UCIHAR) [11]: Detect human activity based on three-axial linear acceleration and three-axial angular velocity that has been captured at a constant rate of 50 Hz. The training and testing data sets are taken from the Human Activity Recognition data set. This data set contains 10,299 samples, each with 561 attributes.

Image Recognition (IMAGE) [12]: Recognize hand-written digits 0 through 9. The training and testing data sets are taken from the Pen-Based Recognition of Handwritten Digits data set. This data set consists of 44 subjects writing each numerical digit 250 times. The samples from 30 subjects are used for training and the other 14 are used for testing.

Accuracy

In this section, we study the impact of quantization and two-level search classification accuracy. Table 1 shows the HD classification accuracy for four different configurations when we categorize the class hypervectors into groups of $m = 1$ to 4. The configuration $m = 1$ is the baseline HD, where we do not have any grouping. HD in $m = 2$ configuration is where each group consists of two class hypervectors. This generates k/m hypervectors in the category stage. Our evaluation shows that grouping has a minor impact on classification accuracy (0.6% on average). HD classification accuracy is also a weak function of grouping configurations. However, the number of hypervectors in the category stage affects the number of computations needed

for inference. Therefore, we choose the grouping approach that minimizes the number of required operations for a given application. For example, for activity recognition with $k = 12$ classes, the grouping with $m = 4$ results in maximum efficiency since it reduces the number of effective hypervectors from $k = 12$ to $k/m + m = 7$.

Table 1 also shows the HD classification accuracy for two types of quantization. Our results show that HD on an average loses 3.7% in accuracy when quantizing the trained model values to the power of two values (2^i). However, quantizing the values to $2^i + 2^j$ values enables HD to provide similar accuracy to HD with integers with less than 0.5% error. This quantization results in $2.2\times$ energy efficiency improvement and $1.6\times$ speedup by modeling the multiplication with two shifts and a single add operation.

Efficiency

The goal is to have HD to be small and scalable so that it can be stored and processed on embedded devices with limited resources. In the conventional HD, each class is represented using a single hypervector. We address this issue by grouping classes together, which significantly lowers the number of computations, and with quantization, which removes costly multiplications from the similarity check.

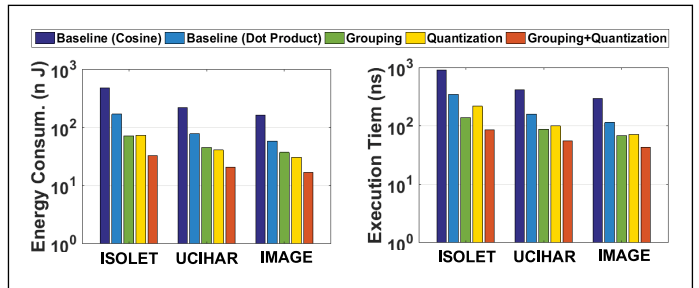


Figure 3. Energy consumption and execution time of HD using proposed optimization approaches.

Figure 3 compares the energy consumption and execution time of the proposed approach with the baseline HD computing during inference. We reported results such that the reader can see the impact of different optimizations. To have a fair comparison, the baseline HD uses the same encoding and number of retraining iterations as the proposed design. Our evaluation shows that grouping of class hypervectors can achieve on an average 5.3× energy efficiency improvement and 4.9× faster as compared to the baseline HD using cosine similarity. In addition, quantization ($2^i + 2^j$) of class elements can further improve the HD efficiency by removing costly multiplications. Our evaluations show that HD enhancing with both grouping and quantization achieves 11.6× energy efficiency and 8.3× speedup as compared to baseline HD using cosine while providing similar classification accuracy.

HD COMPUTING IS a promising solution to performing lightweight classification tasks; however, HD is computationally expensive when working with applications that have large numbers of classes. In this article, we proposed three novel approaches to reduce the HD computation cost during inference. Our first approach simplifies cosine similarity operations to dot product operations by caching class hypervector magnitudes. Our second approach reduces the number of computations during inference by grouping the class hypervectors and performing the similarity check in two stages. Our third approach quantizes the HD trained model and removes costly multiplications from the similarity check. Using the proposed approach enables us to exploit HD as a lightweight classifier for computing on the edge, such as on small embedded devices. ■

Acknowledgments

This work was supported in part by CRISP, one of six centers in JUMP, an SRC program sponsored by DARPA, and in part by the NSF under Grant 730158, Grant 1527034, and Grant CNS-1339335.

References

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*. Lake Tahoe, NV, USA: NeurIPS, pp. 1097–1105, 2012.
- [2] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cogn. Comput.*, vol. 1, no. 2, pp. 139–159, 2009.
- [3] A. Rahimi, P. Kanerva, and J. M. Rabaey, "A robust and energy-efficient classifier using brain-inspired hyperdimensional computing," in *Proc. 2016 Int. Symp. Low Power Electr. Des.*, pp. 64–69.
- [4] M. Imani et al., "Low-power sparse hyperdimensional encoder for language recognition," *IEEE Des. Test*, vol. 34, no. 6, pp. 94–101, 2017.
- [5] A. Rahimi et al., "Hyperdimensional computing for blind and one-shot classification of EEG error-related potentials," *Mobile Netw. Appl.*, vol. 1, pp. 1–12, 2017.
- [6] M. Imani et al., "A binary learning framework for hyperdimensional computing," in *Proc. DATE*, 2019.
- [7] M. Imani et al., "HDcluster: An accurate clustering using brain-inspired high-dimensional computing," in *Proc. DATE*, 2019.
- [8] M. Imani et al., "Hierarchical hyperdimensional computing for energy efficient classification," in *Proc. 55th Annu. Des. Autom. Conf.*, p. 108, 2018.
- [9] M. Imani et al., "VoiceHD: Hyperdimensional computing for efficient speech recognition," in *Proc. Int. Conf. Rebooting Computing (ICRC)*, pp. 1–6, 2017.
- [10] "UCI machine learning repository." [Online]. Available: <http://archive.ics.uci.edu/ml/datasets/ISOLET>
- [11] D. Anguita et al., "Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine," in *Proc. Int. Workshop Ambient Assisted Living*, Springer, pp. 216–223, 2012.
- [12] "UCI machine learning repository." [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/Pen-Based+Recognition+of+Handwritten+Digits>

Mohsen Imani is currently pursuing the PhD degree with the Department of Computer Science and Engineering, University of California San Diego, CA. His research interests include brain-inspired computing and computer architecture.

Justin Morris is currently pursuing the PhD degree with the Electrical and Computer Engineering Department, University of California San Diego, CA. Morris has a BS from the Department of Computer Science and Engineering, University of California San Diego (2018). His research interests include hyperdimensional computing, machine learning, and processing in-memory.

Helen Shu is currently pursuing the BSc degree with the Computer Science and Engineering Department, University of California San Diego, CA. Her research interests include hyperdimensional computing and machine learning.

Shou Li is a graduate student with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA. Li has a BS from the Department of Computer Science and Engineering, University of California San Diego, CA (2018). Her research interest includes machine learning.

Tajana Rosing is a Professor, a holder of the Frattamico Endowed Chair, and the Director of the System Energy Efficiency Laboratory at University of California San Diego, CA. Her research interests include energy efficient computing and embedded and distributed systems.

■ Direct questions and comments about this article to Mohsen Imani, University of California San Diego, La Jolla, CA 92093 USA; moimani@ucsd.edu.